arXiv:1603.04556v4 [math.ST] 25 Feb 2017

Estimation in Tournaments and Graphs under Monotonicity Constraints

Sabyasachi Chatterjee and Sumit Mukherjee

5734 S. University Avenue
Chicago, IL 60637
e-mail: sabyasachi@galton.uchicago.edu

1255 Amsterdam Avenue
New York, NY 10027
e-mail: sm3949@columbia.edu

Abstract: We consider the problem of estimating the probability matrix governing a tournament or linkage in graphs from incomplete observations, under the assumption that the probability matrix satisfies natural monotonicity constraints after being permuted in both rows and columns by some latent permutation. This condition is classical in the social sciences literature (see Fishburn (1973)) and has been studied in the statistics literature in recent work (see, Chatterjee (2015), Shah et al. (2016)). In this paper, we investigate in detail a natural estimator which bypasses the need to search over all possible latent permutations and hence is computationally tractable, and derive asymptotic risk bounds for our estimator. In addition, we prove an automatic adaptation property of our estimator for several sub classes of our parameter space which are of natural interest. These sub classes include generalizations of the popular Bradley Terry Model in the Tournament case, the β model and Stochastic Block Model in the Graph case, and Hölder matrices in the tournament and graph settings.

1. Introduction

In this paper we consider two statistical estimation problems. We begin by describing the two set ups.

- Consider the situation of n teams playing in a league tournament where each team plays every other team once. The results of the tournament can be written as a data matrix y of zeroes and ones by setting $y_{ij} = 1$ for $i < j$ if team i wins against team j, and 0 otherwise. Let θ_{ij} be the probability that team i wins against team j with $\theta_{ji} = 1 - \theta_{ij}$ whenever $i \neq j$. Set $\theta_{ii} = 0$ for all $1 \leq i \leq n$ as a matter of convention. The upper triangular part of the data matrix y is modeled as

$$y_{ij} \sim Bern(\theta_{ij}), \ \ \forall 1 \leq i \leq j \leq n \tag{1.1}$$

where y_{ij} in the upper triangular part is jointly independent and $Bern(.)$ refers to the standard Bernoulli distribution.The lower triangular part of the data matrix is filled in an anti-symmetric manner; that is

$$y_{ij} = 1 - y_{ji}, \ \ \forall 1 \leq j < i \leq n. \tag{1.2}$$

The problem then is to estimate the pairwise comparison probability matrix θ based on the observed data matrix y. This setting can arise whenever the data is in the form of pairwise comparisons (see David (1963)), for example in analyzing customer preferences for items, citation patterns for journals (see Stigler (1994)). For convenience, we stick to the tournament terminology in this paper. Note that we have $O(n^2)$ parameters to estimate and $O(n^2)$ data points. Hence one needs structural assumptions on θ for consistent estimation to be possible. The classical approach in this problem is to assume that the pairwise probability matrix has the following structural form:

$$\theta_{ij} = \frac{\exp(w_i)}{\exp(w_i) + \exp(w_j)}. \tag{1.3}$$

where the $w = (w_1, \ldots, w_n)$ vector is a vector of weights representing the skill/ability of the teams. This is the Bradley-Terry model (see Bradley and Terry (1952)) which is very popular in the ranking literature. It is then common to estimate w by maximum likelihood (see Hunter (2004)) and plug it in to estimate θ_{ij}. The study of asymptotic estimation of w in the Bradley Terry Model has a long history in Statistics (see Simons et al. (1999) and references therein).

We are interested in the problem of estimating the matrix of probabilities θ_{ij} under an assumption commonly made in the ranking literature known as Strong Stochastic Transitivity (**SST**) (see Shah et al. (2016) and references therein). This assumption posits the existence of an ordering among the teams which is unknown to the statistician. This ordering is then reflected on the probabilities θ_{ij} as follows. Let team j have a higher rank than team k (i.e. team j is better than team k). Then for any team i, the probability of team i defeating team k would be no less than the probability of team i defeating team j, which gives $\theta_{ij} \leq \theta_{ik}$.

Even though the SST condition is classical, a formal statistical study under this condition was formally introduced very recently in Chatterjee (2015) and the model described above was termed as the *Nonparametric Bradley Terry Model*. The terminology is apt because it clearly generalizes the very commonly used Bradley Terry model. Any matrix θ of the Bradley Terry form in (1.3) satisfies the SST condition with the ordering given by the ordering of the w vector. We refer to Shah et al. (2016) who show, in a precise sense, that the Nonparametric Bradley Terry Model is a significant extension of the usual Bradley Terry model.

In many realistic scenarios, we would not be able to observe all pairwise comparisons. Thus, it is of interest whether one can still estimate the pairwise comparison matrix θ in the situation where we observe only a fraction of all possible games that can be observed. In this

paper we consider the missing data at random setting. In this setting we get to observe each entry above the diagonal with probability p independently of other observations above the diagonal.
The purpose of this paper is to propose and analyze a computationally tractable estimator in this problem. The main focus of this paper is to obtain finite sample risk bounds (upto a constant factor) and study how small can p be to still allow consistent estimation.

- Consider now the situation of observing a random graph on n nodes with no self loops. Let θ_{ij} now be the probability of node i and node j being linked. Again we set $\theta_{ii} = 0$ for all $1 \leq i \leq n$ as a matter of convention. The random graph can be now encoded as an adjacency matrix y of zeroes and ones. Again, the upper triangular part of the adjacency matrix is modelled as

$$y_{ij} \sim Bern(\theta_{ij}) \ \forall \ 1 \leq i \leq j \leq n \tag{1.4}$$

where y_{ij} in the upper triangular part is jointly independent. The lower triangular part of the data matrix is now filled in a symmetric manner; that is

$$y_{ij} = y_{ji} \ \forall \ 1 \leq j < i \leq n. \tag{1.5}$$

Inspired by the SST assumption in the ranking literature, here we assume that that the vertices can be arranged in an order (unknown to the statistician) of increasing tendency of getting linked to other vertices. This assumption will again impose monotonicity constraints on the edge probabilities θ_{ij}. For example if node j is more "active" or "popular" than node k then for any node i we must have $\theta_{ik} \leq \theta_{ij}$. For an example where such an assumption seems natural, consider a social network with n people labeled $\{1, 2, \cdots, n\}$ where the i^{th} person has a popularity parameter $p_i \in [0, 1]$. The chance that person i and person j are friends is $f(p_i, p_j)$, where f is increasing in both co-ordinates to signify that increasing popularity leads to more friendship ties. The function f also needs to be symmetric, as the chance that i and j are friends is symmetric in (i, j). Indeed, in this case there is (at least) one ordering which sorts the nodes of the network in increasing order of popularity.
We pose and study the problem of estimating the edge probability matrix θ in this set up with missing data at random. We sometimes refer to this model of random graphs as the symmetric model, differentiating it from the antisymmetric (tournament) case. Under our model assumptions, the problem of estimating the edge probabilities is very closely related to the problem of estimating graphons in the spirit of Gao et al. (2015) where we assume monotonicity (without smoothness) of the graphon in both variables, instead of smoothness assumptions made in Gao et al. (2015).

In this paper we look at the above two estimation problems in the antisymmetric and the symmetric model in a unified way. In particular, the purpose of this paper is to introduce and study the risk properties of a natural estimator which is described in subsection 1.2. Our estimator has the

same form in both the models and the technique of analyzing the risk properties of the estimator in both the models is the same.

1.1. Formal Setup of our problem

In this subsection we define two parameter spaces; one for the antisymmetric model and one for the symmetric model.

Denote S_n to be the set of all permutations on n symbols. For any $n \times n$ matrix θ and any permutation $\pi \in S_n$ we define $\theta \circ \pi$ to be the $n \times n$ matrix such that $(\theta \circ \pi)_{ij} = \theta_{\pi(i),\pi(j)}$. If Π denotes the $n \times n$ permutation matrix corresponding to the permutation $\pi \in S_n$, then we have $\theta \circ \pi = \Pi^T \theta \, \Pi$, where the RHS is usual matrix multiplication.

Let $\mathcal{T}$ be the space of tournament matrices defined by

$$\mathcal{T} := \{\theta \in [0,1]^{n \times n} : \theta_{ij} \leq \theta_{ik} \; \forall i < k < j; \quad \theta_{ji} = 1 - \theta_{ij} \; \forall \; i \neq j; \quad \theta_{ii} = 0 \; \forall i\}. \tag{1.6}$$

Any matrix in $\mathcal{T}$ when only looked at the upper triangular part above the diagonal is non increasing in any row (as j grows) and non decreasing in any column (as i grows). The lower triangular part is just 1 minus the upper triangular part and the diagonals are zero. In words, $\mathcal{T}$ is the space of matrices which satisfy the SST assumption with known ranking where the ranking is such that player n is the best, followed by player $n-1$ and so on. Then our parameter space for the antisymmetric model can be written as

$$\Theta_{\mathcal{T}} = \{\theta \circ \pi : \theta \in \mathcal{T}, \pi \in S_n\}. \tag{1.7}$$

Similarly, define the space of matrices

$$\mathcal{G} = \{\theta \in \mathbb{R}^{n \times n} : \theta_{ij} \leq \theta_{ik} \; \forall i < j < k; \quad \theta_{ji} = \theta_{ij} \; \forall i \neq j; \quad \theta_{ii} = 0 \; \forall i\}. \tag{1.8}$$

Any matrix in $\mathcal{G}$ when only looked at the upper triangular part above the diagonal is non decreasing in both rows and columns. The lower triangular part is symmetrically filled, and the diagonals are zero. Again, $\mathcal{G}$ is the space of expected adjacency matrices which are consistent with the monotonicity restrictions imposed by the ordering where node n is most popular followed by node $n-1$ and so on. Then our parameter space for the symmetric model can be written as

$$\Theta_{\mathcal{G}} = \{\theta \circ \pi : \theta \in \mathcal{G}, \pi \in S_n\}. \tag{1.9}$$

For $\Theta = \Theta_{\mathcal{T}}$ or $\Theta_{\mathcal{G}}$ we study the problem of estimating the underlying matrix of probabilities θ. The loss function we consider is the mean Frobenius squared metric defined for any two matrices θ and $\widetilde{\theta}$ as $\frac{1}{n^2}\|\widetilde{\theta} - \theta\|^2$, where $\|A\|$ denotes the Frobenius norm of the matrix A.

1.2. The estimator

The purpose of this paper is a detailed study of the following estimator which consists of several steps. Initially our data matrix y will have some entries 1 or 0 and some entries would be missing.

Let $\widehat{p}$ denote the proportion of non missing entries in our data. That is, $\widehat{p}$ equals the number of observed games divided by $\binom{n}{2}$.

(a) **Zero Filling**
Fill the missing entries of the data matrix y by 0 above the diagonal and fill antisymmetrically below the diagonal. Henceforth let us denote y to be the data matrix filled with zeroes on the upper diagonal wherever we have a missing entry.

(b) **Sorting**
Let $r_i = \sum_{j=1}^{n} y_{ij}$ be the i^{th} row sum of the data matrix y. In this step we sort the vertices according to the row sums $(r_1, \ldots, r_n)$ of the data matrix y and obtain a permutation $\widehat{\sigma}$ such that $r_{\widehat{\sigma}(1)} \leq \cdots \leq r_{\widehat{\sigma}(n)}$. The random permutation $\widehat{\sigma}$ can be thought of as an estimator of π^{-1} which is the inverse of the latent permutation π. We call this step the sorting step, which we use to construct a sorted data matrix $y \circ \widehat{\sigma}$ defined as $(y \circ \widehat{\sigma})_{ij} = (y_{\widehat{\sigma}(i)}, y_{\widehat{\sigma}(j)})$. In case there are ties in the vector $r = (r_1, \ldots, r_n)$, break the ties uniformly at random while obtaining the permutation $\widehat{\sigma}$.

(c) **Scaling and Projection**
In this step we divide the sorted data matrix by $\widehat{p}$. We then project $\frac{1}{\widehat{p}}(y \circ \widehat{\sigma})$ onto the relevant parameter space. In the antisymmetric model, we project onto the set $\mathcal{T} \subset [0,1]^{n \times n}$ and in the symmetric model, we project onto the set $\mathcal{G} \subset [0,1]^{n \times n}$. Both $\mathcal{T}$ and $\mathcal{G}$ are closed convex sets of matrices and hence there exists a unique projection onto them. Let the projection operator be denoted by $Proj$ in both cases. After this step we have the projection of a scaled and sorted data matrix $Proj\Big(\frac{1}{\widehat{p}}(y \circ \widehat{\sigma})\Big)$.

(d) **Unsorting**
We now unsort the scaled, projected and sorted data matrix $Proj\Big(\frac{1}{\widehat{p}}(y \circ \widehat{\sigma})\Big)$ by applying to it the permutation $\widehat{\sigma}^{-1}$.

(e) **Zeroing out**
We now define our final estimator $\widehat{\theta}$ as follows:

$$\begin{aligned} \widehat{\theta} :=& Proj\Big(\frac{1}{\widehat{p}}(y \circ \widehat{\sigma})\Big) \circ \widehat{\sigma}^{-1} \text{ if } \widehat{p} \geq \frac{1}{n}, \\ =& 0 \text{ otherwise.} \end{aligned} \tag{1.10}$$

Remark 1.1. Note that in the antisymmetric (tournament) model, the row sums of the data matrix y correspond to the number of wins or victories for each player, and the column sums of the data matrix correspond to the number of defeats for each player. Hence our sorting step just sorts the teams according to the number of victories (or equivalently the number of defeats, as sum of victory and defeat of each player is $n-1$). Similarly, in the symmetric (graph) model, the row sums of the adjacency matrix y correspond to the empirical degrees of the nodes in the graph. Therefore, our sorting step sorts the vertices according to the empirical degrees.

Remark 1.2. The first step of our algorithm just needs computation of the row sums and then sorting, which combined clearly will only take at most $O(n^2)$ operations. The projection step is thus

going to be dominating the computation time. Fortunately there are efficient ways of computing the projection. First of all, the projection by definition has to be zero on the diagonals and would be anti symmetric/symmetric according to our model. Hence in either of the models, it suffices to compute the projection for the upper diagonal part. It turns out that the spaces $\mathcal{T}$ and $\mathcal{G}$, without the constraint that all elements have to live in $[0,1]$, can be viewed as the space of Isotonic functions on an appropriate Directed Acyclic Graph (DAG) on the domain $\{(i,j) : 1 \leq i < j \leq n\}$. Isotonic functions on a DAG can only increase on following a directed path in the DAG. It can be checked that the problem of computing the projection $Proj$ can be done by first performing Isotonic regression on the appropriate DAG and then hard threshholding the values to lie between 0 and 1. Recent results in Kyng et al. (2015) study how to compute such Isotonic projections on a general DAG. Their result imply an algorithm with runtime $O(n^3)$ for our problem. More classically, the problem of computing this projection is closely related to computing what is called a Bivariate Isotonic Regression(see Chatterjee et al. (2016)). There exists efficient iterative algorithms for this purpose; see page 27 in Robertson et al. (1988). The focus of this paper therefore is not on computation of our estimator but rather on its statistical properties.

2. Main result

Having defined our estimator, we now make a couple of more definitions.

Definition 2.1. Henceforth we will use the notation Θ in place of $\Theta_{\mathcal{T}}$ or $\Theta_{\mathcal{G}}$. The implication is that all results hold with Θ replaced by either of the two parameter sets.

Setting $[n] := \{1, 2, \cdots, n\}$, for any $\theta \in \Theta$ define the row sums $R_i(\theta) := \sum_{j \in [n]} \theta_{ij}$ for $i \in [n]$. Also for each $\theta \in \Theta$ and $p \in [0,1]$ define the quantity

$$Q(p,\theta) := \sum_{i \in [n]} \max_{j : |R_j(\theta) - R_i(\theta)| \leq 2\sqrt{2np \log n}} \sum_{k=1}^{n} (p\,\theta_{ik} - p\,\theta_{jk})^2.$$

The quantity $Q(p,\theta)$ is an important quantity in our analysis. As will be clear from our subsequent analysis, matrices in Θ with additional structure tend to have smaller $Q(p,\theta)$. We are now ready to state the main theorem of this paper.

Theorem 2.1. *Consider the estimator $\widehat{\theta}$ defined in* (1.10)*. Let $n \geq 2$ and let $p \geq \frac{2}{n}$. Then there exists a universal constant $C < \infty$ such that for any $\theta^* \in \Theta$ we have*

$$\frac{1}{n^2}\mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\Big[\frac{(\log n)^2}{np} + \frac{Q(p,\theta^*)}{n^2 p^2}\Big]. \tag{2.1}$$

Remark 2.1. Theorem 2.1 gives an upper bound which is a sum of two terms. The first term scaling like $\widetilde{O}(1/np)$ can be thought of as the risk arising due to the shape constraint imposed by the SST condition. The second term involving $Q(p,\theta^*)$ can be interpreted as the risk arising because of our sorting step; it measures how much the θ changes in a frobenius norm squared sense, if its rows/columns are permuted by a typical sorting permutation.

Remark 2.2. The strength of Theorem 2.1 is that it is adaptive in the parameter θ^*, and gives tight asymptotic bounds for several sub-parameter spaces of interest. The term $\frac{Q(p,\theta^*)}{n^2p^2}$ dominates the $\frac{(\log n)^2}{np}$ term in most cases of interest. Hence the quantity $Q(p,\theta^*)$ determines the rate of convergence of our estimator. All our results will be derived as corollaries of our main theorem. We now begin to discuss the various consequences and implications of Theorem 2.1.

2.0.1. ***Worst case Risk Bound***

As a first application of Theorem 2.1, we deduce the worst case risk of our estimator.

Corollary 2.1. *There is a universal constant $C < \infty$ such that for any $p \geq \frac{2}{n}$ and $n \geq 2$ we have*

$$\sup_{\theta^* \in \Theta} \frac{1}{n^2}\mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\Big[\frac{(\log n)^2}{np} + \sqrt{\frac{\log n}{np}}\Big]. \tag{2.2}$$

Remark 2.3. If there is no missing data (i.e. $p = 1$), the brute force LSE in our problem achieves a MSE of $O(n^{-1})$ upto log factors which is minimax rate optimal in the tournament setting (see Shah et al. (2016)). The same fact can be shown to be true in the graph setting by an application of Lemma 3.1 in Chatterjee and Lafferty (2017). We do not carry this out in this manuscript.

Comparing Theorem 2.1 to the minimax rate, the $\widetilde{O}(n^{-\frac{1}{2}})$ rate achieved by our estimator is clearly worse than the minimax rate of estimation. However the only method known to achieve the minimax rate is the LSE, which is perhaps not computationally feasible. This raises the important question of whether there exists a computationally feasible estimator achieving the minimax rate in this problem. This question deserves further study and is beyond the scope of the current manuscript.

The authors in Shah et al. (2016) improved the analysis of an estimator based on singular value threshholding proposed in Chatterjee (2015) and demonstrated its rate of convergence to be $\widetilde{O}(n^{-\frac{1}{2}})$. Hence our estimator matches the best known rate of convergence for computationally feasible estimators.

Remark 2.4. It follows from Corollary 2.1 that the MSE converges to 0 as soon as $np \gg (\log n)^2$. On the other hand, if np converges to a finite number λ, it is not hard to show that no estimator is consistent, i.e. the MSE stays bounded away from 0.

A natural question is whether the upper bound for our estimator in Corollary 2.1 is tight. The following example suggests that the given upper bound is tight at least when $p = 1$.
Define $\theta^{wc} \in \mathcal{T}$ for any $i \neq j$ as

$$\theta^{wc}_{i,j} = \frac{1}{2} + \frac{1}{4}\mathcal{I}\{i > j\}. \tag{2.3}$$

Simulations of the MSE (see figure 2.0.1) suggest that our exponent of n in the upper bound in Corollary 2.1 is tight, atleast when $p = 1$. For the plot we have simulated our data when the true

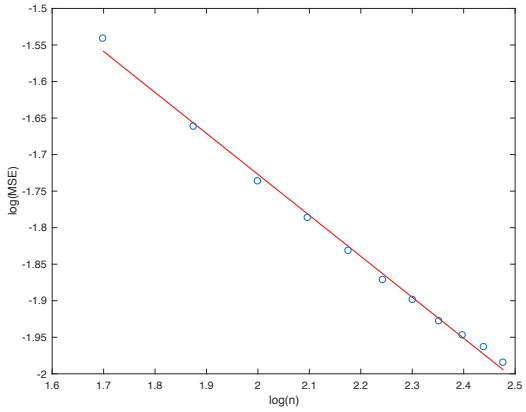

FIGURE 1. *MSE for $\theta = \theta^{wc}$ (worst case behavior)*

underlying $n \times n$ parameter matrix is given by (2.3). We have simulated 100 times for each sample size n ranging from 50 to 300 in increments of 25. A plot of log(n) (base 10) versus log(MSE) clearly shows a linear behavior, and the best fitted line has slope -0.5607, which is close to the predicted exponent -0.5.

2.1. Adaptive Risk Bounds

The main reason for studying our estimator is that it exhibits automatic adaptation properties. By automatic adaptation we mean that even though our estimator achieves $\widetilde{O}(n^{-\frac{1}{2}})$ rate of estimation globally, it achieves provably faster rates of estimation for several subclasses of our parameter space which are of independent interest. We now describe some of these subclasses along with the associated risk bounds for our estimator in these subsets of our parameter space.

2.1.1. Block matrices

Let us consider the class of matrices within our parameter space Θ which are a permuted version of some $k \times k$ blockmatrix. This class of matrices (without the restriction of being within Θ) is popularly known as the Stochastic Blockmodel in Network Analysis. In the Tournament setting, one can think of the situation where there are n teams but only $k \ll n$ levels of teams. Within each level, each team has the same ability. Therefore, assuming a block structure on the true pairwise comparison matrix is a natural way to impose sparsity in the model.

Definition 2.2. For $k \in [n]$, let $\Theta^{(k)} \subseteq \Theta$ denote the subset of all $k \times k$ block matrices with equal sized blocks, upto an unknown permutation.

Corollary 2.2. *There exists a universal constant $C < \infty$ such that for any $p \geq \frac{2}{n}$ we have*

$$\sup_{\theta^* \in \Theta^{(k)}} \frac{1}{n^2} \mathbb{E} \|\widehat{\theta} - \theta^*\|^2 \leq C \Big[\frac{\min(k, \sqrt{np}) \log n}{np} + \frac{(\log n)^2}{np} \Big].$$

for some universal constant $C < \infty$.

Remark 2.5. In particular when k is fixed and $p = 1$, we get from corollary 2.2 that

$$\frac{1}{n^2}\mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\frac{(\log n)^2}{n}.$$

The next theorem shows that this is the minimax error rate for 2×2 blocks of equal size, upto $\log$ factors for our parameter space Θ. In the case when the latent ranking is known, the estimation problem becomes very closely related to Bivariate Isotonic Regression (see Chatterjee et al. (2016)), where the minimax error rate for 2×2 block matrices (without a latent permutation) is shown to be $\widetilde{O}(n^{-2})$. This means that the latent permutation in our parameter space makes the statistical estimation problem harder in a minimax sense.

Theorem 2.2. *Consider the subset $\Theta^{(2)} \subset \Theta$ of all equal sized 2×2 block matrices. Then for some universal constant $c > 0$ we have*

$$\inf_{\widetilde{\theta}} \sup_{\theta \in \Theta^{(2)}} \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n}\mathbb{E}(\widetilde{\theta}_{ij} - \theta_{ij})^2 \geq \frac{c}{n}.$$

Remark 2.6. Without the monotonicity constraint, it was shown in Gao et al. (2015) that for the symmetric $k \times k$ block matrices the minimax rate is $O(\frac{k^2}{n^2} + \frac{\log k}{n})$. Thus for $k \in [1, \sqrt{n}]$ the minimax rate with and without monotonicity constraints are both $\widetilde{O}(n^{-1})$. For $k \in [\sqrt{n}, n]$ the minimax rate with monotonicity is $\widetilde{O}(n^{-1})$, whereas without monotonicity the minimax rate is strictly larger.

2.1.2. Generalized Bradley Terry model + Generalized β model

In the Tournament setting, if the Bradley Terry model is correct, one can use the MLE of the weights vector w to estimate θ. It is known that this estimator attains a fast rate of convergence of $O(n^{-1})$ (see Theorem 4 in Shah et al. (2016)). A natural question is whether our estimator attains this fast rate of convergence whenever the true matrix is actually of the Bradley Terry form. Our next corollary shows this to be true in a more general sense. Before stating the corollary, let us define the Generalized Bradley Terry Model $\Theta_{gbt,M}$ contained within the class of SST matrices.

Definition 2.3. Let $\Theta_{gbt,M} \subset \Theta_{\mathcal{T}}$ denote the subset of all tournament matrices such that $\theta_{i,j} = F(w_i - w_j)$, where F is a unknown symmetric distribution function on $\mathbb{R}$ (i.e. $F(x) + F(-x) \equiv 1$) with a continuous strictly positive density function function, and $\{w_i\}_{i\in[n]}$ is a sequence of real numbers in $[-M, M]$. It is not hard to check that the class $\Theta_{gbt,M} \subset \Theta_{\mathcal{T}}$ indeed satisfies the SST condition. In particular, setting $F(x) = \frac{e^x}{1+e^x}$ and $F(x) = \Phi(x)$ (the normal distribution function) we get the usual Bradley Terry model and the Thurstone model (see Bradley and Terry (1952); Thurstone (1927)) respectively. Let us call this class of matrices the Generalized Bradley Terry Model.

Let us now define the analogous class of matrices in the symmetric model on graphs.

Definition 2.4. Let $\Theta_{gbm,M} \subset \Theta_{\mathcal{G}}$ denote the subset of all expected adjacency matrices such that $\theta_{i,j} = F(w_i + w_j)$, where F is a distribution function on $\mathbb{R}$ with a continuous strictly positive density function, and $\{w_i\}_{i\in[n]}$ is a sequence of real numbers in $[-M, M]$. In particular for the choice $F(x) = \frac{e^x}{1+e^x}$ gives the β model on networks, which has originated from Social Sciences and has been studied in Statistics (c.f. Chatterjee et al. (2011) and references there-in). The class $\Theta_{gbm,F,M}$ generalizes the usual β model to allow for a more general class of distribution functions.

With F assumed to be known, estimation in the Generalized Bradley Terry model has been studied by Shah et al. (2016). We present our next corollary treating F as an unknown parameter.

Corollary 2.3. *There exists a universal constant $C < \infty$ such that for any $p \geq \frac{2}{n}$ we have*

$$\sup_{\theta^* \in \Theta_{gbt,M}} \mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\Big[\Big(\frac{L}{L'}\Big)^2 \frac{\log n}{np} + \frac{(\log n)^2}{np}\Big],$$

where $L := \sup_{|x|\leq 2M} f(x)$ and $L' := \inf_{|x|\leq 2M} f(x)$. The same risk bound holds for the Generalized Beta Model $\Theta_{gbm,M}$.

Remark 2.7. In particular if $p = 1$ then corollary 2.3 shows that

$$\mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\frac{(\log n)^2}{n}$$

when θ is of the Generalized Bradley Terry form. If the distribution function F is known, the authors in Shah et al. (2016) show that the estimator based on the MLE of w achieves the rate $O(n^{-1})$, under the stronger condition that the density f is strongly log concave and twice differentiable. Obviously, computing the MLE takes into account the explicit knowledge of F. The estimator studied in this paper attains the same fast rate of convergence as the MLE upto log factors, but without the knowledge of F.

Remark 2.8. The minimax rate for the Generalized Bradley Terry Model $\Theta_{gbt,M}$ with F known (satisfying some regularity conditions) and $p = 1$ has been established to be $O(1/n)$, see Theorem 4 in Shah et al. (2016)). This implies that our estimator is minimax rate optimal upto log factors for $\Theta_{gbt,M}$ with F unknown in the $p = 1$ case.

2.1.3. Smooth matrices

In the symmetric model (graph setting) there has been a tradition of studying estimation of graphons satisfying some smoothness condition (see Gao et al. (2015) and references therein). It is also a natural question to ask what happens when the true pairwise probability matrix in the Tournament setting is a smooth matrix, upto an unknown permutation. This motivates us to define the following class of matrices satisfying a discrete version of the usual Holder smoothness conditions.

Definition 2.5. For $\alpha, L > 0$, let $\Theta(\alpha, L) \subset \Theta$ denote the subset of all permuted versions of Holder continuous matrices with order α and Holder constant L , i.e. θ, after being permuted in rows and columns, satisfies

$$|\theta_{i,j} - \theta_{k,l}| \leq L\Big(\Big|\frac{i-k}{n}\Big|^{\alpha} + \Big|\frac{j-l}{n}\Big|^{\alpha}\Big)$$

for all $i \neq j, k \neq l, \in [n]$.

The next corollary shows that our estimator provably attains faster rates of convergence than $\widetilde{O}(n^{-\frac{1}{2}})$ whenever the true matrix, upto an unknown permutation, satisfies smoothness conditions as above.

Corollary 2.4. *There exists a universal constant $C < \infty$ such that for any $p \geq \frac{2}{n}$ we have*

$$\sup_{\theta^* \in \Theta(\alpha, L)} \frac{1}{n^2}\mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\Big[\max(1, L)\frac{\log n}{(np)^{\frac{2\alpha+1}{2\alpha+2}}} + \frac{(\log n)^2}{np}\Big]$$

Remark 2.9. As the above corollary shows, the adaptation of our estimator depends crucially on the order of the Holder class α. In particular, for the class of Lipschitz matrices (which correspond to the choice $\alpha = 1$) with no missing entries (which correspond to $p = 1$), the corollary implies the following upper bound:

$$\sup_{\theta^* \in \Theta(1,1)} \mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\,\frac{\log n}{n^{3/4}}.$$

A natural question is whether the worst case MSE over $\Theta(1, 1)$ actually scales like $\widetilde{O}(n^{-3/4})$. We do not know the answer to this question. However, we do have an explicit example of $\theta^* \in \Theta(1, 1)$ for which the quantity $Q(1, \theta^*)$ actually scales like $n^{\frac{5}{4}}$, which suggests that it is not possible to prove a better rate than $\widetilde{O}(n^{-3/4})$ with our proof technique.

Even though the worst case risk over the class of Lipschitz matrices is $\widetilde{O}(n^{-3/4})$, under an extra assumption that θ is lower Lipschitz as well, we get an improved MSE of $\widetilde{O}(n^{-1})$. More generally, the same holds for lower Holder continuous matrices as well. To make this precise we propose the following definition:

Definition 2.6. For $\alpha, L, L' > 0$ let $\Theta(\alpha, L, L') \subset \Theta(\alpha, L)$ denote the subset of all lower Holder continuous matrices with lower Holder constant L', i.e. after being permuted in rows and columns by some permutation it satisfies

$$|\theta_{i,j} - \theta_{k,l}| \geq L'\Big(\Big|\frac{i-k}{n}\Big|^{\alpha} + \Big|\frac{j-l}{n}\Big|^{\alpha}\Big).$$

Corollary 2.5. *There exists a universal constant $C < \infty$ such that for any $p \geq \frac{2}{n}$ we have*

$$\sup_{\theta^* \in \Theta(\alpha, L, L')} \mathbb{E}\|\widehat{\theta} - \theta^*\|^2 \leq C\Big[\Big(\frac{L}{L'}\Big)^2\frac{\log n}{np} + \frac{(\log n)^2}{np}\Big].$$

Remark 2.10. A similar estimator (consists of a sorting step and a smoothing step) as ours was proposed in Chan and Airoldi (2014) in the symmetric model, where the authors assume that the true θ is both lower and upper Lipschitz (but not necessarily monotonic). Under an extra assumption on the sparsity of the gradient of the histogram of θ, (Chan and Airoldi, 2014, Theorem 3) shows that a similar two stage estimator has MSE scaling like $\widetilde{O}(n^{-1})$. We obtain the same error bound without the sparsity assumption, but under the extra assumption of monotonicity. Another point worth mentioning here is that monotonicity constraints make it possible for our estimator to be completely tuning parameter free while the estimator proposed in Chan and Airoldi (2014) has a bandwidth parameter that needs to be tuned.

2.2. Main Contribution

In this paper we study a natural estimator for estimating the matrix of probabilities for two very related models, the tournament model and the graph model under the assumption that the underlying true ordering is unknown with monotonicity constraints. Our estimator matches the best known rate of convergence for computationally feasible estimators. More pertinently, we also demonstrate automatic adaptive properties of our tuning parameter free estimate for interesting sub-classes of the parameter space, which includes the stochastic block model, and the parametric Bradley Terry model and classes of matrices satisfying certain smoothness conditions. Throughout the paper we work with incomplete observations missing at random.

2.3. Scope of future Work

An important open question in the $p = 1$ case is whether there exists a polynomial time estimator achieving the minimax rate of estimation of $\widetilde{O}(n^{-1})$. Another natural question is whether the upper bounds for our estimator given in Corollary 2.4 for Lipschitz matrices and other smooth matrices are tight. Finally, automatic adaptation properties of the Singular Value Threshholding estimator (proposed in Chatterjee (2015) and studied in Shah et al. (2016)) also deserve attention.

3. Proof of corollaries

Proof of Corollary 2.1. To begin, fix $i, j \in [n]$ such that $|R_i(p\theta^*) - R_j(p\theta^*)| \leq 2\sqrt{2np\log n}$. Then we have

$$\sum_{k\in[n]}^{n} [p\theta^*_{ik} - p\theta^*_{jk}]^2 \leq p \sum_{k\in[n]}^{n} |p\theta^*_{ik} - \theta^*_{jk}| = p|R_i(p\theta^*) - R_j(p\theta^*)| \leq 2p\sqrt{2np\log n}.$$

The above bound implies $Q(p, \theta^*) \leq 2\sqrt{2}(np)^{3/2}\sqrt{\log n}$, which along with Theorem 2.1 completes the proof of the corollary. □

Proof of Corollary 2.2. As before, it suffices to control the terms in $Q(p, \theta^*)$. To this end let B be the underlying $k \times k$ matrix, i.e.

$$\theta^*_{i,j} = B\Big(\Big\lceil \frac{ki}{n} \Big\rceil, \Big\lceil \frac{kj}{n} \Big\rceil\Big).$$

Now fix $i, j \in [n]$ such that $|R_i(p\,\theta^*) - R_j(p\,\theta^*)| \le 2\sqrt{2np\log n}$. Then with $s := \Big\lceil \frac{ki}{n} \Big\rceil$ and $t := \Big\lceil \frac{kj}{n} \Big\rceil$ we have

$$2\sqrt{2np\log n} \ge |R_i(p\,\theta^*) - R_j(p\,\theta^*)| = \frac{np}{k}\Big|\sum_{r\in[k]}(B(s,r) - B(t,r))\Big|,$$

which in particular means

$$\max_{r\in[k]} |B(s,r) - B(t,r)| \le \min\Big(1, 2k\sqrt{\frac{2\log n}{np}}\Big).$$

This gives

$$\begin{aligned}
\|p\,\theta^*_{i.} - p\theta^*_{j.}\|^2 =& \frac{np^2}{k}\sum_{r\in[k]}(B(s,r) - B(t,r))^2 \\
&\le p \max_{r\in[k]} |B(s,r) - B(t,r)| \times \frac{np}{k}\Big|\sum_{r\in[k]}(B(s,r) - B(t,r))\Big| \\
&\le p \min\Big(1, 2k\sqrt{\frac{2\log n}{np}}\Big) \times 2\sqrt{2np\log n} \\
&= p \min\Big(2\sqrt{2np\log n}, 8k\log n\Big) \le 8p\min(k, \sqrt{np})\log n.
\end{aligned}$$

The required bound then follows from Theorem 2.1. □

Proof of corollary 2.3. To begin, note that

$$L'|w_i - w_j| \le |\theta^*_{i,k} - \theta^*_{j,k}| \le L|w_i - w_j|.$$

Now fixing $i, j \in [n]$ such that $|R_i(p\,\theta^*) - R_j(p\,\theta^*)| \le 2\sqrt{2np\log n}$, we have

$$2\sqrt{2np\log n} \ge p\sum_{k\in[n]} |\theta^*_{i,k} - \theta^*_{j,k}| \ge npL'|w_i - w_j|,$$

and so

$$\sum_{k\in[n]} [p\,\theta^*_{i,k} - p\,\theta^*_{j,k}]^2 \le np^2L^2(w_i - w_j)^2 \le 8p\Big(\frac{L}{L'}\Big)^2 \log n,$$

which along with Theorem 2.1 completes the proof. □

Proof of Corollary 2.4. As before it suffices to control the terms in $Q(p, \theta^*)$. To this end, fix $i, j \in [n]$ such that $|R_i(p\,\theta^*) - R_j(p\,\theta^*)| \le 2\sqrt{2np\log n}$. We claim that

$$\max_{k\in[n]} |\theta^*_{i,k} - \theta^*_{j,k}| \le \Big(2\sqrt{\frac{2\log n}{np}}\Big)^{\frac{\alpha}{\alpha+1}} (4L)^{\frac{1}{\alpha+1}}. \tag{3.1}$$

We first complete the proof of the corollary, deferring the proof of (3.1). To this end we have

$$\begin{aligned}
\sum_{k\in[n]} [p\,\theta^*_{i,k} - p\,\theta^*_{j,k}]^2 &\le p \max_{k\in[n]} |\theta^*_{i,k} - \theta^*_{j,k}| \times \sum_{k\in[n]} |p\,\theta^*_{i,k} - p\,\theta^*_{j,k}| \\
&\le p\Big(2\sqrt{\frac{2\log n}{np}}\Big)^{\frac{\alpha}{\alpha+1}} (4L)^{\frac{1}{\alpha+1}} \times 2\sqrt{2np\log n} \\
&= p\, 2^{\frac{3\alpha+4}{2\alpha+2}}\, L^{\frac{1}{\alpha+1}}\, (np)^{1-\frac{2\alpha+1}{2\alpha+2}} (\log n)^{\frac{2\alpha+1}{2\alpha+2}} \\
&\le 4p\, \max(1, L)\, (np)^{1-\frac{2\alpha+1}{2\alpha+2}}\, \log n,
\end{aligned}$$

from which the result follows on using Theorem 2.1.

It thus remains to complete the proof of (3.1). To this end, let $\delta := \max_{k\in[n]} |\theta^*_{i,k} - \theta^*_{j,k}|$, and let $k^\star \in [n]$ be such that $|\theta^*_{i,k^\star} - \theta^*_{j,k^\star}| = \delta$. Then setting $N := \Big(\frac{\delta}{4L}\Big)^{1/\alpha} n$, for any $k \in [n]$ such that $|k - k^\star| \le N$ an application of triangle inequality gives

$$\begin{aligned}
|\theta^*_{i,k} - \theta^*_{j,k}| \ge& |\theta^*_{i,k^\star} - \theta^*_{j,k^\star}| - |\theta^*_{i,k} - \theta^*_{i,k^\star}| - |\theta^*_{j,k} - \theta^*_{j,k^\star}| \\
\ge& \delta - \frac{2LN^\alpha}{n^\alpha} = \frac{\delta}{2}.
\end{aligned}$$

This in turn implies

$$2\sqrt{2np\log n} \ge |R_i(p\,\theta^*) - R_j(p\,\theta^*)| \ge p \sum_{k:|k-k^\star|\le N} |\theta^*_{i,k} - \theta^*_{j,k}| \ge pN\delta = p\,\Big(\frac{\delta}{4L}\Big)^{1/\alpha} n\delta,$$

which is same as

$$\delta \le \Big(2\sqrt{\frac{2\log n}{np}}\Big)^{\frac{\alpha}{\alpha+1}} (4L)^{\frac{1}{\alpha+1}}.$$

□

Proof of Corollary 2.5. As usual, fix $i, j \in [n]$ such that $|R_i(p\,\theta^*) - R_j(p\,\theta^*)| \le 2\sqrt{2np\log n}$. Then we have

$$2\sqrt{2np\log n} \ge |R_i(p\,\theta^*) - R_j(p\,\theta^*)| \ge npL'\Big(\frac{'|i-j|}{n}\Big)^\alpha,$$

which implies

$$\sum_{k\in[n]} [p\,\theta^*_{i,k} - p\,\theta^*_{j,k}]^2 \le np^2L^2\Big(\frac{|i-j|}{n}\Big)^{2\alpha} \le 8p\Big(\frac{L}{L'}\Big)^2 \log n,$$

from which the result follows using Theorem 2.1. □

4. Proof of Theorem 2.1

The goal of this section is to prove Theorem 2.1. We accomplish this in several steps.

4.1. Step 1

Recall that our estimator is of the form

$$\widehat{\theta} = Proj\Big(\frac{1}{\widehat{p}} y \circ \widehat{\sigma}\Big) \circ \widehat{\sigma}^{-1}.$$

on the set $n\widehat{p} \geq 1$. We now show that it is sufficient for our purposes to analyse the estimator where $\widehat{p}$ is replaced by p. We now state a lemma making this precise.

Lemma 4.1. *Suppose*

$$\widetilde{\theta} := Proj\Big(\frac{1}{p} y \circ \widehat{\sigma}\Big) \circ \widehat{\sigma}^{-1}$$

be the estimator constructed assuming p known. If $p \geq \frac{1}{n}$, then there exists a constant $C < \infty$ such that for all $n \geq 1$ we have

$$\mathbb{E}\|\widehat{\theta} - \widetilde{\theta}\|^2 \leq \frac{C}{n^2 p^2}.$$

The proof of Lemma 4.1 is given in the appendix. From now on, we will analyze the estimator $\widetilde{\theta}$.

4.2. Step 2

The observed data matrix y is distributed as $Bernoulli(p\theta^* \circ \pi^*)$ where $\theta^* \in \mathcal{T}$ and $\pi^* \in S_n$ are true unknown parameters. Recall that this means on the upper diagonal, $y_{ij} \sim Bernoulli(p\theta^*_{\pi^*(i),\pi^*(j)})$ and they are jointly independent. The lower diagonal is defined to be 1 minus the upper diagonal part. Our task is to estimate $\theta^* \circ \pi^*$. Recall that

$$\widetilde{\theta} = Proj\Big(\frac{1}{p} y \circ \widehat{\sigma}\Big) \circ \widehat{\sigma}^{-1}.$$

Our first step is to write the following inequality.

$$\begin{aligned}
\|\widetilde{\theta} - \theta^* \circ \pi^*\| =& \|Proj\Big(\frac{1}{p} y \circ \widehat{\sigma}\Big) \circ \widehat{\sigma}^{-1} - \theta^* \circ \widehat{\sigma}^{-1} + \theta^* \circ \widehat{\sigma}^{-1} - \theta^* \circ \pi^*\| \\
\leq& \|Proj\Big(\frac{1}{p} y \circ \widehat{\sigma}\Big) \circ \widehat{\sigma}^{-1} - \theta^* \circ \widehat{\sigma}^{-1}\| + \|\theta^* \circ \widehat{\sigma}^{-1} - \theta^* \circ \pi^*\| \\
=& \|Proj\Big(\frac{1}{p} y \circ \widehat{\sigma}\Big) - \theta^*\| + \|\theta^* \circ \widehat{\sigma}^{-1} - \theta^* \circ \pi^*\|.
\end{aligned}$$

The last inequality is just an application of the triangle inequality. We claim that the distribution of both the terms in the last line of the above inequality does not depend on π^*. This is the content of Lemma 6.2 stated and proved in the appendix in Section 6.

By virtue of Lemma 6.2, without loss of any generality we will henceforth assume that the true underlying latent permutation π^* is the identity permutation and therefore $y \sim Bern(p\theta^*)$.

4.3. Step 3

Our starting point in this step is the inequality in the last display of the previous step where we set π^* to be the identity permutation.

$$\|\widetilde{\theta} - \theta^*\| \leq \|Proj\Big(\frac{1}{p}y \circ \widehat{\sigma}\Big) - \theta^*\| + \|\theta^* \circ \widehat{\sigma}^{-1} - \theta^*\|$$

In this step of the proof, we will handle the first term in the right side of the above inequality. Our next lemma concerns this term $\|Proj(\frac{1}{p}y \circ \widehat{\sigma}) - \theta^*\|$. Before stating our next lemma we define the following supremum of an empirical process term which plays an important role in our analysis. Fix $\theta^* \in \mathcal{T} \subset [0,1]^{n \times n}$. Define a function $f_{\theta^*} : \mathbb{R}_+ \to \mathbb{R}$ as follows:

$$f_{\theta^*}(t) = \sup_{\theta \in \mathcal{T} \subset [0,1]^{n \times n} : \|\theta - \theta^*\| \leq t} \langle \frac{1}{p}[y \circ \widehat{\sigma}] - \theta^*, \theta - \theta^* \rangle$$

Here the $\langle\rangle$ notation refers to the inner product of two matrices which are vectors in $[0,1]^{n^2}$ but the inner product is only taken with respect to the entries above the diagonal. The following lemma connects the term $\|Proj\Big(\frac{1}{p}y \circ \widehat{\sigma}\Big) - \theta^*\|$ to f_{θ^*}.

Lemma 4.2. *Fix $\theta^* \in \mathcal{T} \subset [0,1]^{n \times n}$. Let $s > 0$ be a number satisfying the inequality*

$$f_{\theta^*}(s) \leq \frac{s^2}{2}.$$

Then we have the upper bound

$$\|Proj\Big(\frac{1}{p}y \circ \widehat{\sigma}\Big) - \theta^*\| \leq s.$$

Remark 4.1. Lemma 4.2 is a key lemma for us. It is proved using a representation result for $\|Proj(\frac{1}{p}y \circ \widehat{\sigma}) - \theta^*\|$ developed in Chatterjee (2014). Note that this is a deterministic result and does not depend on the distributional properties of y. For a statement and proof of a general version of the above lemma see Lemma 3.1 in Chatterjee and Lafferty (2017). The lemma there is stated for projection onto sets which are a finite union of convex sets whereas we need the result for a convex set.

4.4. Step 4

Lemma 4.2 reduces the problem of upper bounding $\|Proj(\frac{1}{p}y \circ \widehat{\sigma}) - \theta^*\|$ to the problem of finding a $s > 0$ such that $f_{\theta^*}(s) \leq \frac{s^2}{2}$ which in turn behooves us to find a good upper bound of $f_{\theta^*}(t)$ for any fixed t. We now proceed to do this.

Note that $\frac{1}{p}y \circ \widehat{\sigma} = \theta^* \circ \widehat{\sigma} + \frac{1}{p}(y - p\theta^*) \circ \widehat{\sigma}$. This implies the following:

$$f_{\theta^*}(t) \leq \sup_{\theta \in \mathcal{T} : \|\theta - \theta^*\| \leq t} \langle [\theta^* \circ \widehat{\sigma}] - \theta^*, \theta - \theta^* \rangle + \sup_{\theta \in \mathcal{T} : \|\theta - \theta^*\| \leq t} \langle \frac{1}{p}(y - p\theta^*) \circ \widehat{\sigma}, \theta - \theta^* \rangle.$$

An application of the Cauchy Schwarz Inequality to the first term now gets us the following:

$$f_{\theta^*}(t) \leq t\|[\theta^* \circ \widehat{\sigma}] - \theta^*\| + \sup_{\theta \in \mathcal{T}: \|\theta - \theta^*\| \leq t} \langle \frac{1}{p}(y - p\theta^*) \circ \widehat{\sigma}, \theta - \theta^* \rangle. \tag{4.1}$$

The second term on the right side of the above inequality is a suprema of an empirical process. We give this random variable a name by defining

$$M_t = \frac{1}{p} \sup_{\theta \in \mathcal{T}: \|\theta - \theta^*\| \leq t} \langle \epsilon \circ \widehat{\sigma}, \theta - \theta^* \rangle.$$

where $\epsilon = y - p\theta^*$ is a mean zero random matrix with independent entries above the diagonal taking values in $[-1, 1]$.

Also define for any permutation $\pi \in S_n$,

$$M_{\pi,t} = \frac{1}{p} \sup_{\theta \in \mathcal{T}: \|\theta - \theta^*\| \leq t} \langle \epsilon \circ \pi, \theta - \theta^* \rangle.$$

To upper bound M_t we will upper bound the supremum over π of $M_{\pi,t}$. Further, to control the supremum over π of $M_{\pi,t}$ we will upper bound the expectation of $M_{\pi,t}$ for each π and then invoke concentration of $M_{\pi,t}$ around its expectation. The following lemma controls the expectation of $M_{\pi,t}$. The proof is done by a chaining argument and is given in the Appendix.

Lemma 4.3. *Fix any $n > 1$. Fix $\theta^* \in \mathcal{T} \subset [0, 1]^{n \times n}$ and $\pi \in S_n$. Then there exists a universal constant C such that*

$$p \ \mathbb{E} M_{\pi,t} \leq Cn(\log n)^2. \tag{4.2}$$

Now we want to invoke concentration for the random variable $M_{\pi,t}$. To do this we use the following Bernstein inequality like tail bound.

$$P(X^* > \mathbb{E}X^* + u) \leq \exp\big(-\frac{u^2}{2 \sup_{v \in \nu} \mathbb{E}(\langle X, v \rangle^2) + 4\mathbb{E}X^* + 3u}\big) \tag{4.3}$$

Here $X^* = \sup_{v \in \nu} \langle X, v \rangle$ where $X = (X_1, \ldots, X_m)$ is a collection of zero mean random variables taking values in $[-1, 1]$ and ν is a collection of vectors in $[-1, 1]^m$. This result is originally due to Klein et al. (2005) and also appears as Lemma 11 in Shah et al. (2016).

We now apply this result to the random variable

$$pM_{\pi,t} = \sup_{\theta \in \mathcal{T}: \|\theta - \theta^*\| \leq t} \langle \epsilon \circ \pi, \theta - \theta^* \rangle.$$

By (4.2) we have $\mathbb{E}\, p\, M_{\pi,t} \leq Cn(\log n)^2$. Also we have

$$\sup_{\theta \in \mathcal{T}: \|\theta - \theta^*\| \leq t} \mathbb{E}(\langle \epsilon \circ \pi, \theta - \theta^* \rangle)^2 = \sup_{\theta \in \mathcal{T}: \|\theta - \theta^*\| \leq t} \mathbb{E} \sum_{i<j} \epsilon^2_{\pi(i),\pi(j)} (\theta - \theta^*)^2_{ij} \leq pt^2.$$

The last inequality is because variance of each entry of ϵ is atmost p. Hence an application of (4.3) now lets us conclude

$$P(pM_{\pi,t} > \mathbb{E}pM_{\pi,t} + u) \leq \exp\Big(-\frac{u^2}{2pt^2 + Cn(\log n)^2 + 3u}\Big)$$

Define $w = u/p$ to rewrite the above tail bound as

$$P(M_{\pi,t} > \mathbb{E}M_{\pi,t} + w) \leq \exp\Big(-\frac{w^2p^2}{2pt^2 + Cn(\log n)^2 + 3wp}\Big).$$

Now set $w := C\frac{n}{p}(\log n)^2$ for an appropriately large constant C and consider t such that $pt^2 = Cn(\log n)^2$. Then there exists $c > 0$ such that

$$P(M_{\pi,t} > \mathbb{E}M_{\pi,t} + w) \leq \exp(-2cwp) \leq \exp(-2cn(\log n)^2).$$

A union bound then gives

$$P(\max_{\pi \in S_n} M_{\pi,t} > \max_{\pi \in S_n} \mathbb{E}M_{\pi,t} + w) \leq n! \exp(-2cn(\log n)^2) \leq \exp(-cn(\log n)^2)$$

Since $M_t \leq \max_{\pi \in S_n} M_{\pi,t}$ we can finally conclude that

$$\mathbb{P}\Big(M_t > C\frac{n}{p}(\log n)^2\Big) \leq \exp(-cn(\log n)^2).$$

Setting $s^2 = C\frac{n}{p}(\log n)^2$ the above tail bound then implies $f_{\theta^*}(s) \leq \frac{s^2}{2}$ with very high probability. Therefore, by (4.1) we have the bound with probability atleast $1 - \exp(-cn(\log n)^2)$,

$$\|Proj\Big(\frac{1}{p}y \circ \widehat{\sigma}\Big) - \theta^*\| \leq C\frac{n}{p}(\log n)^2 + \|\theta^* \circ \widehat{\sigma}^{-1} - \theta^*\|. \tag{4.4}$$

Lemma 4.2 now gives rise to the following proposition. The proposition follows by noting that the high probability bound can be converted into an expectation bound trivially and by using $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$.

Proposition 4.1.

$$\mathbb{E}\frac{1}{n^2}\|\widetilde{\theta} - \theta^*\|^2 \leq C\frac{(\log n)^2}{np} + \frac{4}{(np)^2}\,\mathbb{E}\|(p\,\theta^* \circ \widehat{\sigma}) - p\,\theta^*\|^2.$$

The above proposition upper bounds the risk of our estimator by a sum of two terms. The first term scaling like $(\log n)^2/np$ term is essentially the minimax rate of our problem (upto logarithmic factors) which is attained by the global least squares estimate (c.f. (Shah et al., 2016, Theorem 5)). The second term is the excess risk of our estimator as compared to the risk of the global least squares estimate. The above proposition now reduces the problem to controlling this excess risk term. This term measures how much the sorting permutation $\widehat{\sigma}$ changes θ^* in a mean sum of squares sense.

4.5. Step 3

In this step we investigate the excess risk term $\frac{1}{(np)^2}\|p\,\theta^* \circ \widehat{\sigma} - p\,\theta^*\|^2$. Analyzing this term is one of the key contributions of this paper. We first prove the following lemma concerning the behavior of $\widehat{\sigma}$. Recall that for any $\theta \in \Theta$ we denote $R_i(\theta) := \sum_{j=1}^n \theta_{ij}$ to be the ith row sum of the matrix θ.

Lemma 4.4. *For any $C > 0$ setting $t_i := \max\{\sqrt{4CR_i(p\,\theta^*)\log n}, 4C\log n/3\}$ we have*

$$\mathbb{P}(|R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*)| > \max(t_i + t_{\widehat{\sigma}_i}, 4C\log n\} \textit{ for some } i \in [n]) \le 2n^{1-C}, \tag{4.5}$$

Proof. Define the vector of row sums of the data matrix $r = (r_1, \ldots, r_n)$ where $r_i = \sum_{j=1}^n y_{ij}$ for all $1 \le i \le n$. Setting

$$A_n := \cap_{i=1}^n \{|r_i - R_i(p\,\theta^*)| \le t_i\} \tag{4.6}$$

we will first show that

$$\mathbb{P}(A_n^c) \le 2n^{1-C}. \tag{4.7}$$

With $\sigma_i^2 := \frac{p}{n}\sum_{j=1}^n \theta_{ij} - \frac{p^2}{n}\sum_{j=1}^n \theta_{ij}^2$, and using the inequality $R_i(p\,\theta^*) \ge n\sigma^2$ we have

$$\frac{t_i^2}{2n\sigma_i^2} \ge \frac{2CR_i(p\,\theta^*)\log n}{n\sigma_i^2} \ge 2C\log n, \quad \frac{3t_i^2}{2t_i} = \frac{3t_i}{2} \ge 2C\log n.$$

Thus an application of Bernstein's inequality gives

$$\mathbb{P}(|r_i - R_i(p\,\theta^*)| > t_i) \le 2\exp\Big\{-\frac{t_i^2/2}{n\sigma_i^2 + t_i/3}\Big\} \le 2e^{-C\log n} = 2n^{-C}.$$

A union bound then proves (4.7).

It thus suffices to show that after conditioning on the event A_n we have for all $1 \le i \le n$,

$$|R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*)| \le \max(t_i + t_{\widehat{\sigma}(i)}, 4C\log n).$$

For this, setting $\widehat{\sigma}(i) = j$, we split the proof into the following cases:

- $\widehat{\sigma}(i) < i, \quad R_i(p\,\theta^*) > 4C\log n$
  In this case we will show that

  $$|R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*)| = R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*) \le t_i + t_{\widehat{\sigma}(i)}.$$

  If not, then we have $R_i(p\theta^*) - t_i > R_j(p\theta^*) + t_j$, and so the intervals $[R_i(p\theta^*) - t_i, R_i(p\theta^*) + t_i]$ and $[R_j(p\,\theta^*) - t_j, R_j(p\,\theta^*) + t_j]$ are disjoint, giving $i = \widehat{\sigma}^{-1}(j) < \widehat{\sigma}^{-1}(i)$. Now, for any $k > i$ we have $R_k(p\,\theta^*) - 4C\log n/3 \le R_i(p\,\theta^*) - 4C\log n/3$ by monotonicity. Also we have

  $$\begin{aligned}
  &(R_k(p\,\theta^*) - \sqrt{4CR_k(p\,\theta^*)\log n}) - (R_i(p\,\theta^*) - \sqrt{4CR_i(p\,\theta^*)\log n}) \\
  =&(R_k(p\,\theta^*) - R_i(p\,\theta^*)) - \sqrt{4C\log n}(\sqrt{R_k(p\,\theta^*)} - \sqrt{R_i(p\,\theta^*)}) \\
  =&(\sqrt{R_k(p\,\theta^*)} - \sqrt{R_i(p\,\theta^*)})\Big[\sqrt{R_k(p\,\theta^*)} + \sqrt{R_i(p\,\theta^*)} - \sqrt{4C\log n}\Big] \ge 0,
  \end{aligned}$$

as $R_i(p\,\theta^*) \geq 4C\log n$. Combining this gives

$$\begin{aligned} R_k(p\,\theta^*) - t_k =& \min(R_k(p\,\theta^*) - 4C\log n/3, R_k(p\,\theta^*) - \sqrt{4CR_k(p\,\theta^*)\log n}) \\ \geq& \min(R_i(p\,\theta^*) - 4C\log n/3, R_i(p\,\theta^*) - \sqrt{4CR_i(p\,\theta^*)\log n}) \\ =& R_i(p\,\theta^*) - t_i \geq R_j(p\,\theta^*) + t_j \end{aligned}$$

This implies that the intervals $[R_k(p\,\theta^*) - t_k, R_k(p\,\theta^*) + t_k]$ and $[R_j(p\,\theta^*) - t_j, R_j(p\,\theta^*) + t_j]$ are disjoint, and so $r_k \geq r_j$ which gives $\widehat{\sigma}^{-1}(k) > \widehat{\sigma}^{-1}(j) = i$. Since this holds for every $k > i$, the permutation $\widehat{\sigma}^{-1}$ maps the set $\{i, i+1, \cdots, n\}$ to a subset of $\{i+1, \cdots, n\}$, which is impossible.

- $\widehat{\sigma}(i) < i, \quad R_i(p\,\theta^*) \leq 4C\log n$
  In this case we have

$$|R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*)| = R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*) \leq R_i(p\,\theta^*) \leq 4C\log n.$$

- $\widehat{\sigma}(i) \geq i$
  In this case we will again show that

$$|R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*)| = R_{\widehat{\sigma}(i)}(p\,\theta^*) - R_i(p\,\theta^*) \leq t_i + t_{\widehat{\sigma}(i)}.$$

  If this does not hold, with $j := \widehat{\sigma}(i)$ we have $R_j(p\,\theta^*) - t_j > R_i(p\,\theta^*) + t_i$. Consequently the intervals $[R_i(p\,\theta^*) - t_i, R_i(p\,\theta^*) + t_i]$ and $[R_j(p\,\theta^*) - t_j, R_j(p\,\theta^*) + t_j]$ are disjoint, and so $r_j > r_i$. By construction of $\widehat{\sigma}$ we have $i = \widehat{\sigma}^{-1}(j) > \widehat{\sigma}^{-1}(i)$. Finally for any $k < i$ we have

$$R_k(p\,\theta^*) + t_k \leq R_i(p\,\theta^*) + t_i \leq R_j(p\,\theta^*) - t_j,$$

  and so the intervals $[R_k(p\theta^*) - t_k, R_k(p\theta^*) + t_k]$ and $[R_j(p\theta^*) - t_j, R_j(p\theta^*) + t_j]$ are disjoint as well. This gives $r_k \leq r_j$, and consequently we have $\widehat{\sigma}^{-1}(k) < \sigma^{-1}(j) = i$. Thus the permutation $\widehat{\sigma}^{-1}$ maps the set $\{1, 2, \cdots, i\}$ to a subset of $\{1, 2, \cdots, i-1\}$, a contradiction.

□

Remark 4.2. The ideal case here is when the permutation $\widehat{\sigma}$ is close to the identity permutation. In general though, $\widehat{\sigma}$ need not be close to the identity permutation. For example, when θ^* is a constant matrix, $\widehat{\sigma}$ is a uniformly random permutation. But Lemma 4.1 shows that $R_{\widehat{\sigma}(i)}$ and R_i are always close irrespective of what θ^* is. In particular, $|R_{\widehat{\sigma}} - R_i| = \widetilde{O}_{[}(\sqrt{n})$ even though both $R_{\widehat{\sigma}}$ and R_i could be $O(n)$.

Remark 4.3. We also note here that Lemma 4.4 can be used to derive concentration results for $\widehat{\sigma}$ under the assumption that the row sums $R_i(\theta^*)$ are strictly increasing. In particular, if $\min_{i\in[n]} |R_i(\theta^*) - R_{i-1}(\theta^*)$ is bounded away from 0, then by Lemma 4.4 with high probability we have

$$\max_{i\in[n]} |i - \widehat{\sigma}(i)| = O(\sqrt{n\log n}).$$

Of course, if the row sums are not increasing, no concentration of $\widehat{\sigma}$ towards identity is expected as stated in the previous remark.

Remark 4.4. We also note that both the bounds above are adaptive in terms of sparsity of the underlying graph. For e.g. if the entries of the matrix θ are mostly 0 or small, the row sums $R_i(\theta)$ will be small as well, thus giving a better bound.

Using Lemma 4.4 we can now control the excess risk term $\frac{1}{n^2p^2}\|p\theta^* \circ \widehat{\sigma} - p\theta^*\|^2$. This is done as follows.

Recall the definition of the event A_n in (4.6). Use Lemma (4.5) with $C = 2$ to note that that for all n large we have

$$\mathbb{P}(A_n) \leq 2n^{-1}, \quad A_n := \cap_{i=1}^n \{|R_i(p\,\theta^*) - R_{\widehat{\sigma}(i)}(p\,\theta^*)| \leq 2\sqrt{2np\log n}\}.$$

This gives

$$\begin{aligned}
&\mathbb{E}\sum_{i,j=1}^n [p\,\theta^*_{ij} - p\,\theta^*_{\widehat{\sigma}(i),\widehat{\sigma}(j)}]^2 \\
=&\mathbb{E}\sum_{i,j=1}^n [p\,\theta^*_{ij} - p\,\theta^*_{\widehat{\sigma}(i),\widehat{\sigma}(j)}]^2 1_{A_n^c} + \mathbb{E}\sum_{i,j=1}^n [p\,\theta^*_{ij} - p\,\theta^*_{\widehat{\sigma}(i),\widehat{\sigma}(j)}]^2 1_{A_n} \\
\leq& n^2p^2\mathbb{P}(A_n^c) + \mathbb{E}\sum_{i,j=1}^n [p\,\theta^*_{ij} - p\,\theta^*_{\widehat{\sigma}(i),\widehat{\sigma}(j)}]^2 1_{A_n} \\
\leq& 2np^2 + 2\mathbb{E}\sum_{i,j=1}^n [p\,\theta^*_{ij} - p\,\theta^*_{\widehat{\sigma}(i),j}]^2 1_{A_n} + 2\mathbb{E}\sum_{i,j=1}^n [p\,\theta^*_{\widehat{\sigma}(i),j} - p\,\theta^*_{\widehat{\sigma}(i),\widehat{\sigma}(j)}]^2 1_{A_n} \\
\leq& 2np^2 + 4Q(p,\theta^*).
\end{aligned}$$

The last display along with Proposition 4.1 completes the proof of Theorem 2.1.

5. Acknowledgements

We want to thank Bodhisattva Sen for introducing us to this problem.

References

Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

Stanley H Chan and Edoardo M Airoldi. A consistent histogram estimator for exchangeable graph models. *arXiv preprint arXiv:1402.1888*, 2014.

Sabyasachi Chatterjee and John Lafferty. Adaptive risk bounds in unimodal regression. 2017. To appear in Bernoulli.

Sabyasachi Chatterjee, Adityanand Guntuboyina, and Bodhisattva Sen. On matrix estimation under monotonicity constraints. 2016. To appear in Bernoulli.

Sourav Chatterjee. A new perspective on least squares under convex constraint. *The Annals of Statistics*, 42(6):2340–2381, 2014.

Sourav Chatterjee. Matrix estimation by universal singular value thresholding. *The Annals of Statistics*, 43(1):177–214, 2015.

Sourav Chatterjee, Persi Diaconis, and Allan Sly. Random graphs with a given degree sequence. *The Annals of Applied Probability*, pages 1400–1435, 2011.

Herbert Aron David. *The method of paired comparisons*, volume 12. DTIC Document, 1963.

John Duchi. Lecture notes for statistics 311/electrical engineering 377. 2016.

Peter C Fishburn. Binary choice probabilities: on the varieties of stochastic transitivity. *Journal of Mathematical psychology*, 10(4):327–352, 1973.

Chao Gao, Yu Lu, Harrison H Zhou, et al. Rate-optimal graphon estimation. *The Annals of Statistics*, 43(6):2624–2652, 2015.

Fuchang Gao and Jon A Wellner. Entropy estimate for high-dimensional monotonic functions. *Journal of Multivariate Analysis*, 98(9):1751–1764, 2007.

David R Hunter. Mm algorithms for generalized bradley-terry models. *Annals of Statistics*, pages 384–406, 2004.

Thierry Klein, Emmanuel Rio, et al. Concentration around the mean for maxima of empirical processes. *The Annals of Probability*, 33(3):1060–1077, 2005.

Rasmus Kyng, Anup Rao, and Sushant Sachdeva. Fast, provable algorithms for isotonic regression in all l_p-norms. In *Advances in Neural Information Processing Systems*, pages 2701–2709, 2015.

Tim Robertson, F. T. Wright, and R. L. Dykstra. *Order restricted statistical inference*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., Chichester, 1988. ISBN 0-471-91787-7.

Nihar B Shah, Sivaraman Balakrishnan, Adityanand Guntuboyina, and Martin J Wainwright. Stochastically transitive models for pairwise comparisons: Statistical and computational issues. In *International Conference on Machine Learning*, 2016.

Gordon Simons, Yi-Ching Yao, et al. Asymptotics when the number of parameters tends to infinity in the bradley-terry model for paired comparisons. *The Annals of Statistics*, 27(3):1041–1060, 1999.

Stephen M Stigler. Citation patterns in the journals of statistics and probability. *Statistical Science*, pages 94–108, 1994.

Louis L Thurstone. A law of comparative judgment. *Psychological review*, 34(4):273, 1927.

Sara A. van de Geer. *Applications of empirical process theory*, volume 6 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2000. ISBN 0-521-65002-X.

RR Varshamov. Estimate of the number of signals in error correcting codes. In *Dokl. Akad. Nauk SSSR*, volume 117, pages 739–741, 1957.

6. Appendix

6.1. Proof of Theorem 2.2

We need to use the following version of Gilbert Varshamov coding lemma (Varshamov (1957)). The proof of this lemma is provided in subsection 6.5.

Lemma 6.1. *[Gilbert-Varshamov] Fix any positive integer d. Let*

$$\mathcal{H}_{1/2} = \{v \in \{0,1\}^d : \sum_{i=1}^{d} v_i = d/2\}.$$

There exists a subset $\mathbb{F} \subset \mathcal{H}_{1/2}$ with

$$|\mathbb{F}| \geq \exp(d/32) \tag{6.1}$$

such that for any $w \neq w^{'} \in W$ we have

$$\frac{d}{2} \geq H(w, w^{'}) \geq \frac{d}{8} \tag{6.2}$$

where H refers to the Hamming distance between any two points of the hypercube.

We are now ready to prove Theorem 2.2.

Proof of Theorem 2.2. For any subset $S \subset \{0,1\}^d$, define a matrix $\theta^S \subset \Theta_2$ as follows:

$$\theta_{ij}^S = \begin{cases} 0.5 & \text{if } i \in S, j \in S \\ 0.5 + \frac{c}{\sqrt{n}} & \text{if } i \in S, j \notin S \\ 0.5 + \frac{c}{\sqrt{n}} & \text{if } i \notin S, j \in S \\ 0.5 + 2\frac{c}{\sqrt{n}} & \text{if } i \notin S, j \notin S \end{cases}.$$

Take two different subsets $S, S^{'}$ of $\{0,1\}^d$. Denote by $S \Delta S^{'}$ to be the symmetric difference of the two sets. Note that whenever $i \in S \Delta S^{'}$ and j does not belong to $S \Delta S^{'}$ or vice versa, we have

$$(\theta_{ij} - \theta_{ij}^{'})^2 \geq \frac{c^2}{n}.$$

Therefore we have

$$\sum_{i=1}^{n} \sum_{j=1}^{n} (\theta_{ij} - \theta_{ij}^{'})^2 \geq 2\frac{c^2}{n} |S \Delta S^{'}| |n - S \Delta S^{'}|. \tag{6.3}$$

Now we apply Lemma 6.1 to extract a finite subset $\mathbb{F} \subset \mathcal{H}_{1/2}$ satisfying (6.1) and (6.2). Consider the set of matrices

$$W = \{\theta \in \Theta : \theta = \theta^S \text{ for some } S \in \mathbb{F}.\}$$

Here we are viewing elements in $\{0,1\}^d$ as subsets of $\{1 \leq i \leq d\}$ in the usual way. Hence the cardinality $|S \Delta S^{'}|$ is exactly the Hamming distance between the corresponding binary strings.

Hence we can now apply (6.2) to (6.3) to get the following bound for any $\theta \neq \theta^{'}$ belonging to W:

$$\sum_{i=1}^{n}\sum_{j=1}^{n}(\theta_{ij} - \theta_{ij}^{'})^2 \geq \frac{c^2 n}{8}. \tag{6.4}$$

This implies that W is a packing set of radius $c\sqrt{n/8}$ with cardinality atleast $\exp(n/32)$. We now upper bound Kulback Leibler divergences. Kulback Leibler divergences between two Bernoulli distributions can be written as follows:

$$D(Bern(p), Bern(p^{'})) = \int_{p^{'}}^{p} \frac{p-x}{x(1-x)} dx. \tag{6.5}$$

In case $0.4 \leq p^{'} < p \leq 0.6$ then we have $KL(Bern(p), Bern(p^{'})) \leq 5(p-p^{'})^2$. In our case, all the entries of our matrix are between 0.4 and 0.6 for large enough n. Hence we then have for all sufficiently large n,

$$\max_{\theta \neq \theta^{'} \in W} D(P_{\theta}, P_{\theta^{'}}) \leq 5 \max_{\theta \neq \theta^{'}} \|\theta - \theta^{'}\|^2 \leq 20nc^2. \tag{6.6}$$

The last inequality holds because each entry of $\theta - \theta^{'}$ is bounded in magnitude by $2c/\sqrt{n}$. So now a standard application of Fano's lemma (see Chapter 13 in Duchi (2016)) with W as our packing set and using (6.4) and (6.6), we obtain a minimax lower bound

$$\inf_{\widetilde{\theta}} \sup_{\theta \in \Theta^{(2)}} \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n}\mathbb{E}(\widetilde{\theta}_{ij} - \theta_{ij})^2 \geq \frac{c^2}{8n}\big(1 - 32\frac{20c^2 n + \log 2}{n}\big).$$

Choosing c appropriately small enough small enough finishes the proof of the theorem. Note that we have ignored the diagonals being 0 in our calculations. Since there are only n diagonal elements, it can be checked that the conclusion remains unchanged. □

6.2. Proof of Lemma 4.1

Proof. To begin, note that on the set $\widehat{p} \geq \frac{1}{n}$ we have

$$\|\widehat{\theta} - \widetilde{\theta}\|^2 \leq \Big(\frac{1}{p} - \frac{1}{\widehat{p}}\Big)^2 \sum_{i,j=1}^{n} y_{ij} \leq n^2 \Big(\frac{1}{p} - \frac{1}{\widehat{p}}\Big)^2 \leq n^4.$$

The first inequality is because the projection operator onto a convex set is a Lipschitz map with Lipschitz constant 1, that is for any two matrices x and y we have $\|Proj(x) - Proj(y)\| \leq \|x - y\|$. The last inequality is because of the assumed condition on $\widehat{p}$ and p. Also on the set $\widehat{p} < \frac{1}{n}$, using $\widehat{\theta} = 0$ we have

$$\|\widehat{\theta} - \widetilde{\theta}\|^2 \leq \frac{n^2}{p^2} \leq n^4.$$

Thus we have

$$\mathbb{E}\|\widehat{\theta}-\widetilde{\theta}\|^2 \le n^4\,\mathbb{P}(\widehat{p}<\frac{p}{2})+n^4\,\mathbb{P}(\widehat{p}>2p)+\mathbb{E}\big[\Big(\frac{1}{\widehat{p}}-\frac{1}{p}\Big)^2\sum_{i,j=1}^{n}y_{ij}\mathcal{I}\Big\{\frac{p}{2}\le\widehat{p}\le 2p\Big\}\big] \tag{6.7}$$

where the last term in the RHS of (6.7) can be bounded by

$$\begin{aligned}\frac{4}{p^4}\mathbb{E}\big[(\widehat{p}-p)^2\sum_{i,j=1}^{n}y_{ij}\mathcal{I}\Big\{\frac{p}{2}\le\widehat{p}\le 2p\Big\}\big] \le&\frac{4n^2}{p^4}\mathbb{E}\big[\widehat{p}\,(\widehat{p}-p)^2\mathcal{I}\Big\{\frac{p}{2}\le\widehat{p}\le 2p\Big\}\big]\\ \le&\frac{8n^2}{p^3}\mathbb{E}(\widehat{p}-p)^2\le\frac{16}{p^2}.\end{aligned}$$

For the first two terms in the RHS of (6.7), using the multiplicative form of Chernoff's bound we have

$$\mathbb{P}\Big(\widehat{p}<\frac{p}{2}\Big)\le e^{-\frac{n(n-1)p}{16}}\le e^{-\frac{n-1}{8}},\quad \mathbb{P}\Big(\widehat{p}>2p\Big)\le e^{-\frac{n(n-1)p}{6}}\le e^{-\frac{n-1}{3}}.$$

Thus plugging in the estimates in (6.7) we get

$$\frac{1}{n^2}\mathbb{E}\|\widehat{\theta}-\widetilde{\theta}\|^2\le n^4(e^{-\frac{n-1}{8}}+e^{-\frac{n-1}{3}})+\frac{16}{n^2p^2},$$

from which the result follows. □

6.3. Proof of Lemma 6.2

Lemma 6.2. *Let $\theta^*\in\mathcal{T}$ and $\pi^*\in S_n$. Also let $y\sim Bern(p\,\theta^*\circ\pi^*)$. Then the distribution of $\|Proj(\frac{1}{p}y\circ\widehat{\sigma})-\theta^*\|$ and $\|\theta^*\circ\widehat{\sigma}^{-1}-\theta^*\circ\pi^*\|$ do not depend on π^*.*

Proof. We set $p=1$ without loss of generality. Setting $\mathcal{H}:=\{0,1\}^{n\times n}\cap\mathcal{T}$, let

$$W=\{w\in\mathcal{H}: R_1(w)\le R_2(w)\cdots\le R_n(w)\}$$

denote the space of all matrices in $\mathcal{H}$ with sorted row sums, and note that the random variable $y\circ\widehat{\sigma}$ takes values in W. For any $w\in W$, define

$$T(w)=\{v\in\mathcal{H}: w=v\circ\pi^*\text{ for some }\pi^*\in S_n\}.$$

In words, $T(w)$ is the set of all matrices in $\mathcal{H}$ which on sorting gives w.

For a given matrix $v\in T(w)$ there are multiple ways of obtaining a sorted version if there are ties in the row sums. If there are ties in the row sums of v we break the ties uniformly at random. Let $p_v(w)$ denote the probability of obtaining w as a sorted version under uniform breaking of ties. Then we can write for any $\theta^*\in\mathcal{T}$,

$$P_{\theta^*}(y\circ\widehat{\sigma}=w)=\sum_{v\in T(w)}P_{\theta^*}(y=v)p_v(w)=\sum_{v\in T(w)}P_{\theta^*\circ\pi^*}(y=v\circ\pi^*)p_{v\circ\pi^*}(w),$$

where we use the facts $P_{\theta^*}(y = v) = P_{\theta^* \circ \pi^*}(y = v \circ \pi^*)$ and $p_v(w) = p_{v \circ \pi^*}(w)$, both of which are easy to verify. To verify the second equality one has to use the fact that we randomize uniformly in the case of ties. Finally observe that the set $T(w)$ is closed under permutations, that is, if $v \in T(w)$ then $v \circ \pi^* \in T(w)$ for any $\pi^* \in S_n$, and so

$$\sum_{v \in T(w)} P_{\theta^* \circ \pi^*}(y = v \circ \pi^*) p_{v \circ \pi^*}(w) = \sum_{v \in T(w)} P_{\theta^* \circ \pi^*}(y = v) p_v(w) = P_{\theta^* \circ \pi^*}(y \circ \widehat{\sigma} = w)$$

where we have again used $p_v(w) = p_{v \circ \pi^*}(w)$. Now the last two displays let us conclude that for all $\pi^* \in S_n$,

$$P_{\theta^*}(y \circ \widehat{\sigma} = w) = P_{\theta^* \circ \pi^*}(y \circ \widehat{\sigma} = w)$$

This therefore implies that the distribution of $y \circ \widehat{\sigma}$ is free from π^* and hence the distribution of $\|Proj(y \circ \widehat{\sigma}) - \theta^*\|$ is free from π^*.

Proceeding to verify the second claim, let $\widehat{\sigma}_1$ be a random permutation with the distribution of $\widehat{\sigma}$ when the underlying matrix is θ^*, and let $\widehat{\sigma}_2$ be a random permutation with the distribution of $\widehat{\sigma}$ when the underlying matrix is $\theta^* \circ \pi^*$. Then for any $\tau \in S_n$ we have

$$\mathbb{P}(\widehat{\sigma}_2^{-1} = \tau) = \mathbb{P}(\widehat{\sigma}_1^{-1} \circ \pi^* = \tau),$$

and so $\widehat{\sigma}_2^{-1} \stackrel{d}{=} \widehat{\sigma}_1^{-1} \circ \pi^*$. Thus we have

$$\|\theta^* \circ \widehat{\sigma}_2^{-1} - \theta^* \circ \pi^*\| \stackrel{d}{=} \|\theta^* \circ \widehat{\sigma}_1^{-1} \circ \pi^* - \theta^* \circ \pi^*\| = \|\theta^* \circ \widehat{\sigma}_1^{-1} - \theta^*\|$$

where the first equality refers to equality in distribution. This immediately implies that $\|\theta^* \circ \widehat{\sigma}^{-1} - \theta^* \circ \pi^*\|$ is also free from π^*, thus completing the proof of the lemma.

□

6.4. Proof of Lemma 4.3

In order to prove the above lemma, we will need to make use of the following two standard results. We first set up some notations. For any set $A \subset \mathbb{R}^n$ define its covering number at radius $\epsilon > 0$ to be the minimum number of Euclidean balls of radius ϵ centred inside A such that the union of the balls is a superset of A. Denote this covering number by $N(A, \epsilon)$ where the Euclidean metric used is implicit.

Let us now define the space of matrices which are non decreasing in both rows and columns with entries between 0 and 1 as

$$\mathcal{M}_{[0,1]} = \{\theta \in [0,1]^{n \times n} : \theta_{ij} \leq \theta_{kl} \text{ iff } i \leq k \text{ and } j \leq l\}.$$

Estimates of the metric entropy of $\mathcal{M}_{[0,1]}$ are available in the literature. The next proposition shows that these metric entropy bounds of $\mathcal{M}_{[0,1]}$ can be used to derive a similar bound for the covering number of $\mathcal{T}$.

Proposition 6.1. *Fix a positive integer n. We have the covering number inequality for any $\epsilon > 0$,*

$$\log N(\epsilon, \mathcal{T}, \|.\|) \leq C \left(\frac{n}{\epsilon}\right)^2 \left[\log\left(\frac{n}{\epsilon}\right)\right]^2.$$

The same upper bound for the log covering number holds when $\mathcal{T}$ is replaced by $\mathcal{G}$.

Proof. Let us prove the proposition for the set $\mathcal{T}$ and the corresponding statement for $\mathcal{G}$ will follow similarly. Let us define a map $D : \mathcal{T} \to [0,1]^{n\times n}$ as follows:

$$D(\theta)_{i,j} = \min\{\theta_{i,i-1}, \theta_{i+1,i}\}\, \mathcal{I}\{i = j\} + \theta_{i,j}\, \mathcal{I}\{i \neq j\}$$

Let us define another map $\phi : [0,1]^{n\times n} \to [0,1]^{n\times n}$ as

$$\phi(\theta)_{i,j} = \theta_{i,n-j+1}$$

Define the composition map $f = \phi \circ D$. Recall that $\mathcal{M}_{[0,1]}$ is the space of monotone matrices with entries in $[0,1]$. It can now be checked that $f(\theta) \in \mathcal{M}_{[0,1]}$ for all $\theta \in \mathcal{T}$ and f is a one to one mapping. We now make the observation that for any $\theta \neq \theta'$ belonging to $\mathcal{T}$ we have

$$\|\theta - \theta'\|^2 \leq \|f(\theta) - f(\theta')\|^2. \tag{6.8}$$

Now also note that f is a continuous map and hence its image $f(\mathcal{T})$ is a closed subset of $\mathcal{M}_{[0,1]}$. Hence there exists a covering set F of $f(\mathcal{T}) \subset \mathcal{M}_{[0,1]}$ at radius ϵ with cardinality atmost $N(\epsilon/2, \mathcal{M}_{[0,1]})$. This is because the set of projections of the covering set F at radius $\epsilon/2$ onto the closed set $f(\mathcal{T})$ forms a covering set of $f(\mathcal{T})$ at radius ϵ. Therefore it follows from (6.8) that the inverse image of F under the mapping f is a covering set of $\mathcal{T}$ at radius ϵ. Thus we can conclude,

$$\log N(\epsilon, \mathcal{T}) \leq \log N(\epsilon/2, \mathcal{M}_{[0,1]}).$$

The covering number for $\mathcal{M}_{[0,1]}$ can be obtained using Lemma 3.4 in Chatterjee et al. (2016) and is written below. This result is proved by using the covering number results in Gao and Wellner (2007).

$$\log N(\epsilon, \mathcal{M}_{[0,1]}) \leq C \left(\frac{n}{\epsilon}\right)^2 \left[\log\left(\frac{n}{\epsilon}\right)\right]^2$$

where $\epsilon > 0$ and C is a universal constant. The last two displays finish the proof. The same proof for $\mathcal{G}$ goes through by defining ϕ to be the identity map. □

We are now ready to prove Lemma 4.3.

Proof of Lemma 4.3. Fix $\theta^* \in \mathcal{T}$. Denote $y - p\theta^*$ by ϵ.

For ease of notation, define

$$\mathcal{A} = \{(\theta - \theta^*) \circ \pi : \theta \in \mathcal{T}\}.$$

An application of a standard chaining result (see van de Geer (2000)) results in the upper bound

$$\mathbb{E}M_{\pi,t} \leq \frac{12}{p} \inf_{0<\delta\leq n} \left\{ \int_{\delta}^{n} \sqrt{\log N(\epsilon, \mathcal{A})}\, d\epsilon + 4n\delta \right\}.$$

The upper limit of the integral is n since the diameter (maximum Euclidean pairwise distance) of $\mathcal{A}$ is bounded by n. Setting $\delta = \frac{1}{n}$ in the last equation we then obtain

$$\mathbb{E}M_{\pi,t} \leq \frac{12}{p} \left\{ \int_{\frac{1}{n}}^{n} C\left(\frac{n}{\epsilon}\right) \left[\log\left(\frac{n}{\epsilon}\right)\right] d\epsilon \right\} + 4.$$

Note that $\log\left(\frac{n}{\epsilon}\right) \leq \log n^2$ because the variable of integration ϵ ranges from $1/n$ to t. This then implies the upper bound

$$\mathbb{E}M_{\pi,t} \leq C\frac{n}{p}(\log n)^2.$$

□

6.5. Proof of Lemma 6.1

Proof. Assume 4 divides d for simplicity. Let $H(.,.)$ denote the Hamming distance on $\{0,1\}^d$, i.e. for any $v_1, v_2 \in \{0,1\}^d$ we have

$$H(v_1, v_2) := \sum_{i=1}^{n} |v_1(i) - v_2(i)|.$$

Take $\beta = \frac{1}{2}$ and $\alpha = \frac{1}{8}$, and $\gamma = 32$. With $d = 4k$ and setting $\widetilde{P} := \{w \in \{0,1\}^{2k} : \sum_{i=1}^{2k} w_i = k\}$ we claim that there exists a set $\widetilde{F} \subset \widetilde{P}$ such that $|\widetilde{F}| \geq e^{d/32}$, and for any $w_1, w_2 \in \widetilde{F}$ we have $H(w_1, w_2) \geq \frac{d}{16}$. Given the claim, define F as

$$F := \{v \in \{0,1\}^d : (v_1, \cdots, v_{2k}) \in \widetilde{F}, v_{2k+1} = v_{2k+2} = \cdots = v_{3k} = 1, v_{3k+1} = \cdots = v_{4k} = 0\},$$

and note that for any $v \in F$ we have $\sum_{i=1}^{4k} v_i = 2k$. Also

$$v_1, v_2 \in \widetilde{P} \Rightarrow \frac{d}{16} \leq H(v_1, v_2) \leq 2k = \frac{d}{2}.$$

Since $|F| = |\widetilde{F}| \geq e^{d/32}$, the proof of the lemma is complete.

It thus remains to prove the claim. To this effect, let $\widetilde{F}$ denote the largest packing set of $\widetilde{P}$ of radius $\frac{d}{16}$ in Hamming distance, i.e. for any $w_1, w_2 \in \widetilde{F}$ we have $H(w_1, w_2) \geq \frac{d}{16}$. Then using the maximaility of $\widetilde{F}$ we have

$$\widetilde{P} = \bigcup_{w\in\widetilde{F}} B_H\left(w, \frac{d}{8}\right) \Rightarrow |\widetilde{P}| \leq |\widetilde{F}| \times \sup_{w\in\widetilde{P}} \left|B_H\left(w, \frac{d}{8}\right)\right|,$$

which implies

$$\frac{1}{|\widetilde{F}|} \leq \frac{\sup_{w\in\widetilde{P}} \left|B_H\left(w, \frac{d}{8}\right)\right|}{|\widetilde{P}|} \leq \frac{\sup_{w\in\{0,1\}^d} \left|B_H\left(w, \frac{d}{8}\right)\right|}{2^{2k}} \frac{2^{2k}}{|\widetilde{P}|} = \frac{\sup_{w\in\widetilde{P}} \mathbb{P}\Big(H(w,\mathbf{X}) \leq \frac{d}{8}\Big)}{\mathbb{P}(\sum_{i=1}^{2k} X_i = k)}$$

where $\mathbf{X} = (X_1, \cdots, X_{2k})$ is a $2k$ dimensional random vector with uniform distribution on $\{0,1\}^{2k}$. To bound the numerator in the RHS of (**??**), note that the vector $(|w_1 - X_1|, |w_2 - X_2|, \cdots |w_n - X_n|)$ is also distributed uniformly on $\{0,1\}^{2k}$, and so the above bound can be rewritten as

$$\frac{1}{|\widetilde{F}|} \leq \frac{\mathbb{P}\Big(B \leq \frac{k}{2}\Big)}{\mathbb{P}(B = k)} \tag{6.9}$$

where $B \sim Bin(2k, .5)$.

The numerator in the RHS of (6.9) can be bound using Hoeffding's inequality as

$$\mathbb{P}\Big(B \leq \frac{k}{2}\Big) \leq e^{-2\frac{(k/2)^2}{2k}} = e^{-k/4},$$

whereas for the denominator, Stirling's approximation gives

$$\mathbb{P}(B = k) = \frac{\binom{2k}{k}}{2^{2k}} \sim \frac{1}{\sqrt{\pi k}}.$$

Combining these two estimates and using (6.9), for all large k we have $|\widetilde{F}| \geq e^{k/8}$, thus completing the proof of the lemma.

□